

### ACTC/Research Performance Tools for Scientific Programming Current Status and On-going Work

David Klepacki, Luiz Derose, Anshul Gupta Advanced Computing Technology IBM T.J. Watson Research

klepacki@us.ibm.com





## **ACTC Tools**

#### Goal

- Complement IBM application performance tools offerings
- Methodology
  - Development in-house and in collaboration with universities, research labs, and industry
  - Deployment "AS IS" to make it rapidly available to the user community for use and evaluation
  - Close interaction between tools builders and application developers for quick feedback for functionality enhancements

#### Disclaimer:

- > Member of IBM Research
- Many designs invented in IBM Research
- Not all became product



David Klepacki



### **Disclaimer**

Member of IBM Research
Many designs invented in
IBM Research
Not all became product



David Klepack

klenacki@us ibm con



## **Current Collaboration Efforts**

- CEPBA (UPC) Jesus Labarta
  - > Performance analysis of OpenMP programs
  - > UTE / DPCL / Paraver interface
- U. Illinois Dan Reed, David Padua, & Josep Torrellas
  - > Pablo performance analysis and visualization system & I/O
  - > Compiler support for data distribution and locality
- U. Maryland Jeff Hollingsworth
  - > Dyninst on AIX
  - > Sigma & dynamic system for performance analysis
- **U.** Wisconsin Bart miller
  - > Paradyn on AIX
  - > Tools scalability on large parallel systems

David Klepaci



### **Outline**

- HPM Toolkit
  - > Status and on-going work
- DPCL Tools
  - > CEPBA Collaboration
- UTE Tools
  - ▶ UTE Gantt Chart
- SiGMA
- WSMP
  - Watson Sparse Matrix Package

David Klenack

klepacki@us.ibm.com



# PMAPI (AIX 5)

- Performance monitor API (PMAPI) developed by the IBM server division at the Austin lab
- Provide performance monitor data access at various levels
- Extension of the regular thread context

David Klepack



#### **HPM Toolkit Goals**

- Correlate parallel source code with hardware and software performance measurements
- Help to identify causes of performance problems
- Low instrumentation overhead
- Flexible and easy to use interface
- Multi-language support
- Portability (AIX & Linux)

David Klenack

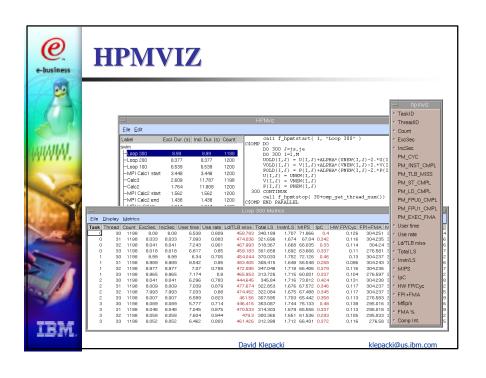
klepacki@us.ibm.con

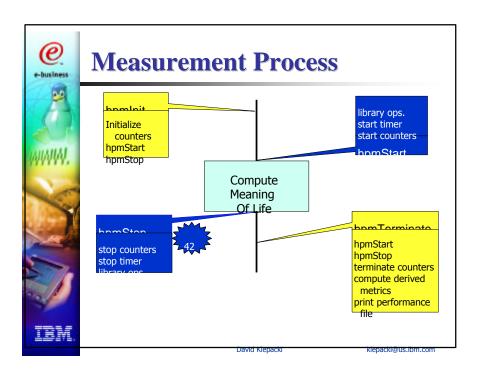


## **HPM Toolkit Components**

- hpmcount
  - > Starts a program and at the end of the execution provides a summary output with:
    - Wall clock time
    - Resource utilization statistics
    - Hardware performance counters information
    - \* Derived hardware metrics
- libhpm
  - Instrumentation library for performance measurement of Fortran, C, and C++ applications
- hpmviz
  - Graphical user interface for visualization of hardware performance data

David Klepacki







### **Current Status**

- Downloadable via http://alphaworks.ibm.com
- Power 4 and AIX 5.1 support
- Automatic hardware counters profiler
  - Function level (loop level with OpenMP)
  - Binary instrumentation with DPCL
  - Functionality to filter or to select functions
  - > Call graph generated during instrumentation
- New Visualizer
  - > Call graph visualization support
  - > Functionality to combine metrics from different runs

David Klenacki

klepacki@us.ibm.com



### **AUTOHPM**

- Automatic HPM instrumentation
- Simple interface
  - > > autohpm a.out
- Support
  - serial
  - > threaded
  - > MPI
  - mixed applications
- Instrumentation
  - > Functions
  - Loops (OpenMP)

David Klepack



### **ACTC - CEPBA Collaboration**

- CEPBA performance tools
  - > Tracing
    - OMPtrace and OMPItrace
  - > Visualization and analysis
    - Paraver

David Klepack

klenacki@us ibm cor

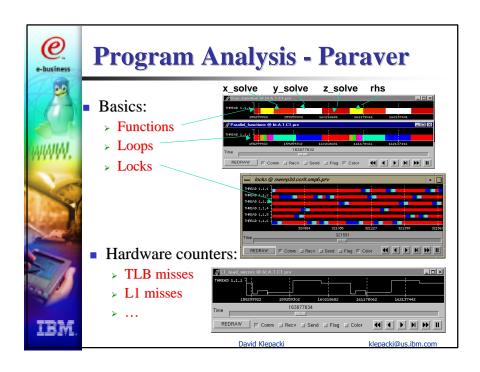


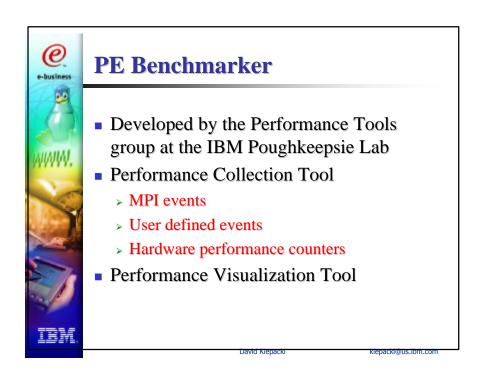
## **OMPtrace and OMPItrace**

- OpenMP and MPI instrumentation
- Interfaces with
  - > DPCL
  - > PMAPI
- Automatic instrumentation
  - Entry and exit of user functions with parallel loops + user specified functions
  - Parallel functions generated by the compiler to encapsulate loop body
  - > MPI calls

David Klepack

klepacki@us.ibm.con



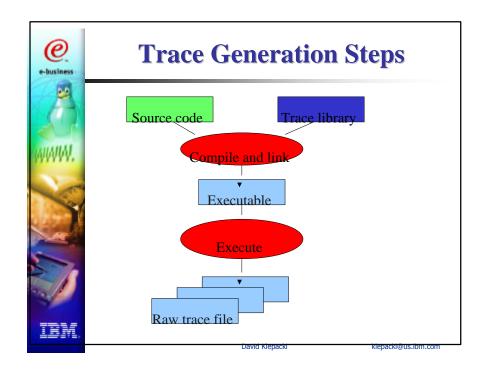


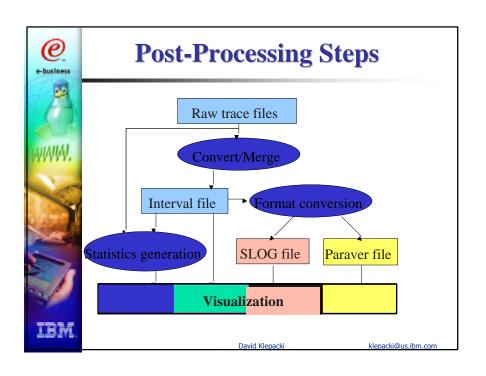


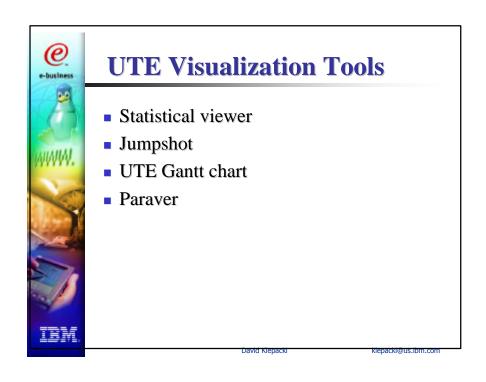
### **Unified Trace Environment - UTE**

- Provides a unified and easily expandable trace environment for MPI and other software layers on IBM SP systems
- Uses AIX trace facility
- Uses standard PMPI profiling interface
- Trace files generated independently
- Generates periodic global clock record
- Provides routines to generate user-defined events

David Klepack







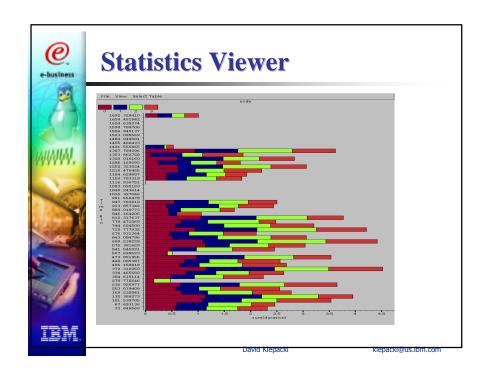


# **Statistics Generation Utility**

- Reads individual UTE interval files or merged ones
- A declarative language. Example: table name=sample condition=(start < 2)</li>
   x=("node", node) x=("processor", cpu)
   y=("avg(duration)", dura, avg)
- Field names (start, node, cpu, duration) are defined in the profile
- Pre-defined statistics

David Klepack

klepacki@us.ibm.cor



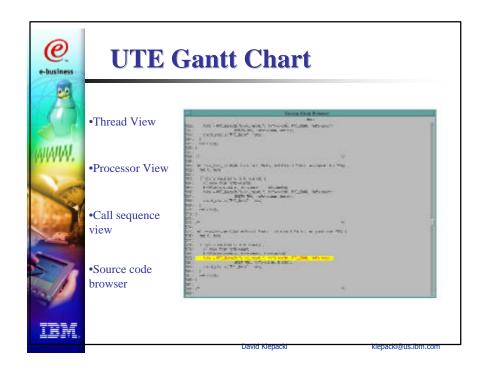


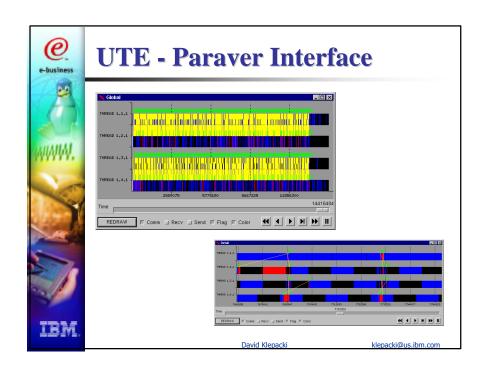
### **UTE Gantt Chart**

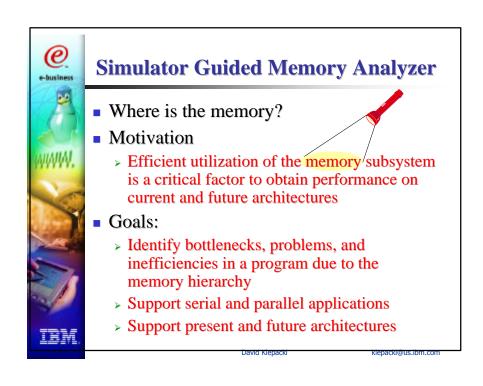
- Reads individual or merged files
- Functions: zoom in/out, hide/show
- Integrated views: many in one
  - Disconnected-state thread view
  - > Nested-state thread view
  - > CPU (processor) view
  - > Thread-processor view
- Call sequence viewer
- Source code browser

David Klepack

klepacki@us.ibm.con









### **SiGMA Overview**

- Family of tools to understand the memory subsystem
  - Focus on detailed statistics
  - > Complement existing hardware counters
- Ability to handle real applications
  - Fortran and C
  - Serial and parallel programs
- Provide hints about restructuring
  - > Padding (both inter and intra data structures)
  - Loop restructuring
  - Blocking parameters

David Klenack

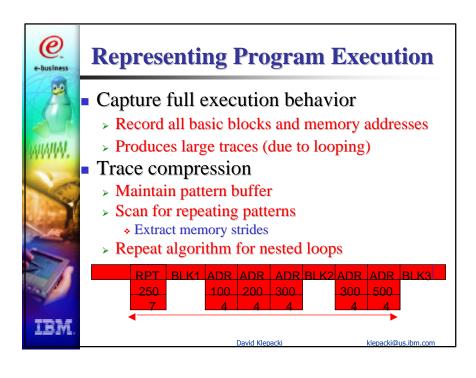
klenacki@us ihm com

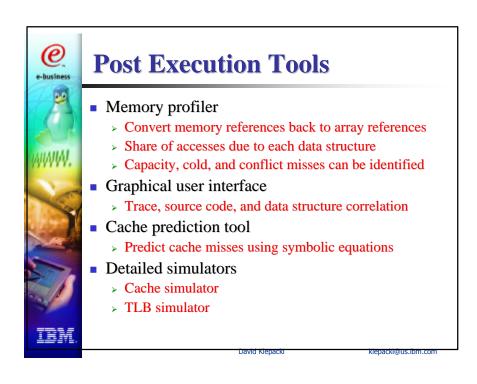


# **Approach**

- Instrument program automatically
- Run instrumented program
  - > Capture full information about memory use
  - > Produce compact trace
    - Extracts loops and memory strides
- Post execution tools

David Klepacki







# **Sigma Memory Dump**

- Takes as input a trace file and the specifications of the memory subsystem.
  - > Supports write-back, write-through caches, and the hardware prefetcher.
- Runs the cache simulation and dumps data into a 4D repository containing:
  - Control-flow (modules, functions, code segments)
  - Data structures
  - Metrics (LoadMisses, LoadHits etc)
  - ➤ Cache levels (L1, L2, TLB ...)

David Klepack

klepacki@us.ibm.cor